

# Ion@XDC2017

Laura Abbott (labbott@redhat.com)

[https://etherpad.openstack.org/p/LPC2017\\_Android\\_Mobile](https://etherpad.openstack.org/p/LPC2017_Android_Mobile)

# A short Ion history

- Memory allocation framework originally written by Rebecca Schultz Zavin at Google for Android
- Designed to manage memory pools
- These days, mostly for getting buffers to gralloc, uses dma\_buf

# Do we still need Ion

- Yes, but in a different way
- Simple wrapper for dma\_buf allocations to userspace (kernel?)
- Useful for gralloc and gralloc like frameworks

# What happened this last year?

- Ripped out a bunch of code and simplified things
- New allocation interface and ABI
  - Just an ioctl for allocation and querying the heaps
- Patches are pending for update to AOSP
  - Thanks John and Sumit!

Why is Ion still in staging - A B |

# Split /dev/ion

- Lots of interest in going from /dev/ion -> /dev/heap1, /dev/heap2
- Patches to review for this
- Question: Discoverability?

# The allocation ABI

```
struct ion_allocation_data {  
    __u64 len;  
    __u32 heap_id_mask;  
    __u32 flags; ← Some people are running out of space?  
    __u32 fd;  
    __u32 unused; ← Okay to expand into here?  
};
```

# Query ioctl

```
struct ion_heap_query {  
    __u32 cnt; /* Total number of heaps to be copied */  
  
    __u32 reserved0; /* align to 64bits */  
  
    __u64 heaps; /* buffer to be populated */  
  
    __u32 reserved1;  
  
    __u32 reserved2;  
  
};
```

# Questions

- Flags field has 16 bits reserved for vendor usage
- Already running out of space(!)
  - Expand? Tell people to be more clever?
- Need to access internal flags for state
  - New ioctl is `_flag_set`, returns bool?

# New heaps, Old heaps

- Addition of new heaps must have at least one open source user
- Delete old heaps for people to bring back later

# Debug accounting

- Memory leaks are painful, how to find out which process is at fault?
  - Someone forgot to close an fd, not surfaceflinger
- Idea: use lsof to show open fds, give dma\_bufs unique names
- Enough correlation?
- Other ideas?

# Testing

- kselftest integration
- Allocation/free easy to test
- In kernel mapping APIs harder to test
- This is roughly a dma-buf test?

# libion

- Currently exists as an AOSP git repo
- Similar to what libdrm provides, wrappers around ioctls
- Okay to just leave in AOSP based on discussion in LPC

# Destaging?

- Happening soon but not too soon
- Live under drivers/gpu and work through drm tree?

Any other discussion?